

# Sequential Bayesian Inference in Hidden Markov Stochastic Kinetic Models with Application to Detection and Response to Seasonal Epidemics

Junjing Lin and Michael Ludkovski

Department of Statistics & Applied Probability, UC Santa Barbara \*

January 17, 2013

## Abstract

We study sequential Bayesian inference in continuous-time stochastic kinetic models with latent factors. Assuming continuous observation of all the reactions, our focus is on joint inference of the unknown reaction rates and the dynamic latent states, modeled as a hidden Markov factor. Using insights from nonlinear filtering of jump Markov processes we develop a novel sequential Monte Carlo algorithm for this purpose. Our approach applies the ideas of particle learning to minimize particle degeneracy and exploit the analytical jump Markov structure. A motivating application of our methods is modeling of seasonal infectious disease outbreaks represented through a compartmental epidemic model. We demonstrate inference in such models with several numerical illustrations and also discuss predictive analysis of epidemic countermeasures using sequential Bayes estimates.

**Keywords:** sequential Monte Carlo, particle learning, jump Markov process, stochastic epidemic models

## 1 Introduction

Stochastic jump-Markov models have become ubiquitous in multiple application areas, including systems biology, molecular chemistry, epidemiology, queuing theory and finance. An important class of such systems is described using the chemical reaction system paradigm, which classifies jumps in terms of a finite number of possible reactions. System transitions are specified probabilistically in terms of the distributions of the inter-reaction periods and next-to-fire reaction type. The reaction rates depend on the current system state and impose a Markovian structure, termed a stochastic kinetic model (SKM) [Wilkinson 2006].

While the basic setup assumes time-stationarity, in many contexts time-dependence, seasonality, and other regime shifts of reaction rates are crucial. A popular way to incorporate stochastic shifts in the environment is through Markov modulation, i.e. introduction of an additional (latent) dynamic factor that affects transition rates. We christen such systems *Hidden Markov Stochastic Kinetic Models* (HMSKM).

Usage of an HMSKM in an application requires statistical estimation of the reaction rates and environmental factors. In the present paper we are concerned with Bayesian inference which allows unified filtering of latent system states and parameters and full quantification of the posterior

---

\*E-mail: linj@pstat.ucsb.edu, ludkovski@pstat.ucsb.edu

uncertainty. Moreover, anticipating dynamic optimization and decision making applications, we are interested in sequential inference. As a motivating example of such HMSKM inference, we describe below a stochastic model of *seasonal epidemics* of infectious diseases. Here the SKM paradigm is used to give a mechanistic description of outbreak progression using a compartmental description of the population, while the latent factor represents environmental factors affecting the outbreak (such as new genetic shifts in the pathogen or weather patterns). Sequential inference of the reaction rates (infectiousness, etc.) and dynamic seasonality determines outbreak severity and is the central problem in biosurveillance and corresponding public health policy-making.

Bayesian inference in SKMs is typically accomplished through Markov chain Monte Carlo (MCMC) methods [Golightly and Wilkinson 2006, Boys et al 2008, Niemi 2009, Golightly and Wilkinson 2011] and has been also used for related epidemics inference problems [O’Neill 2002, Jewell et al 2009, Lawson 2009, Merl et al 2009]. But sequential inference requires re-running an MCMC analysis after every new data point is obtained and therefore is not computationally efficient. A more suitable and flexible alternative is to apply sequential Monte Carlo methods (SMC), also known as particle filters, which use an empirical selection-mutation mechanism [Doucet et al 2001, Cappé et al 2005]. The main drawback of SMC is particle degeneracy which becomes especially serious during estimation of constant parameters. A recent class of *particle learning* methods [Carvalho et al 2011, Dukic et al 2010] has been designed to overcome these challenges by exploiting additional analytical structures. Starting with these ideas we develop new SMC methods targeting continuous-time stochastic kinetic models, demonstrating the efficiency and tractability of Bayesian inference in this novel context.

The aims of this work are three-fold. First, we wish to draw the attention of the computational statistics and stochastic simulation community to the convenient analytic structure of hidden Markov stochastic kinetic models. This structure allows for highly efficient Bayesian sequential inference algorithms in rather general settings, which to our knowledge supercede previous results in this direction. Second, we provide a new extension of the particle learning SMC method. In contrast to existing literature, we work in continuous time and further consider a hidden Markov factor. Thus, to compute the predictive and conditional likelihoods we use tools from stochastic filtering (namely filtering of doubly-stochastic Poisson processes) which have not been hitherto used in this context. Our results demonstrate the wide applicability and attractiveness of particle learning in jump-Markov models.

Finally, we extend the burgeoning body of literature on sequential inference in compartmental epidemiological models. Precisely, we propose a new SIR-type model of seasonal epidemics with a stochastic seasonality factor. The model allows for tractable online inference assuming full observation of all epidemiological events and can be used as a testbed to analyze alternative ways of epidemic control under imperfect information.

In the next section we provide more background on our main motivating application in epidemic modeling; the rest of the paper is organized as follows. Section 2 provides a general setup of a hidden Markov stochastic kinetic model; the resulting inference problem and the particle learning algorithm are constructed in Section 3. Section 4 then illustrates the results and provides numerical examples for a simple SIS model of seasonal epidemics. Section 5 then discusses predictive analysis of outbreak countermeasures in the developed sequential framework.

## 1.1 Seasonal Epidemics

Probabilistic modeling of infectious disease epidemics is an important tool in public health analysis. Stochastic models provide a tractable way to quantify the uncertainty about epidemic dynamics and carry out predictive analysis on the future path of the outbreak. Public health

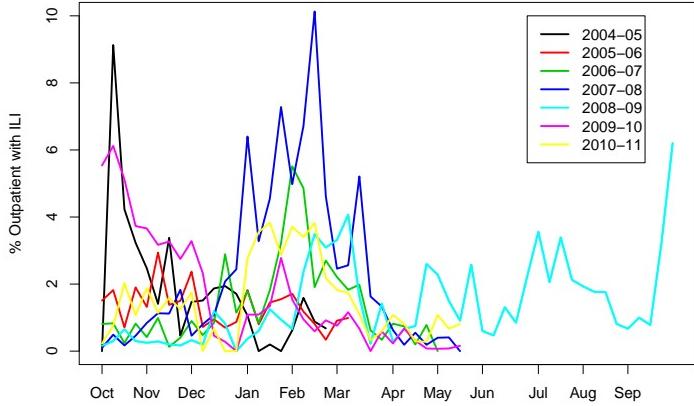


Figure 1: Percentage of Influenza Like Illness (ILI) outpatient visits in Santa Barbara county in 2004-2011 based on weekly sentinel providers data collected by Santa Barbara County Department of Public Health [Bellomy 2011]. No data is normally collected during the summer.

agencies increasingly rely on such techniques to compare alternative courses of action and mount an efficient policy response through behavioral or biomedical interventions.

Recurring seasonal epidemics, with a prime example of influenza, provide some of the best testbeds for quantitative analysis thanks to availability of time-series data of previous outbreaks. Figure 1 shows the ILI flu incidence statistic in Santa Barbara county in California over the past seven years. ILI or influenza-like illness is a formal description of symptoms that typically occur in infected individuals, and is the most common proxy for measuring the incidence of flu in the population. The figure illustrates the apparent fact that flu is generally more prevalent during winters, due to emergence of new virus strains and colder weather making susceptibility higher. Nevertheless, significant variation in flu incidence cases and timing of outbreaks can be observed year over year. A case in point was the 2008-09 season when the new H1N1 influenza strain caused a world-wide pandemic peaking in early Fall 2009, scrambling the traditional flu calendar of public health actions<sup>1</sup>. To summarize, Figure 1 demonstrates three main features of flu outbreaks: (i) strong seasonality; (ii) year-to-year variability; and (iii) stochastic fluctuations during each outbreak. Other endemic infectious diseases with similar patterns include rotavirus, norovirus, measles, dengue fever, and cholera [Grassly and Fraser 2006].

Full understanding of community epidemic dynamics remains elusive given the paucity of available data. While issues such as inhomogeneous mixing, age effects, and spatial interactions are undoubtedly crucial, their mathematical and statistical modeling requires large-scale computational and modeling efforts [Halloran et al 2008]. Alternatively, *mechanistic* models of outbreak progression that can be calibrated to real data. A popular mechanistic approach is given by the class of stochastic compartmental models [Andersson and Britton 2000]. Thus, the population is partitioned into several classes of individuals based on their epidemiological status, such as Susceptible, Infected, etc., and the outbreak is described on the macroscopic level in terms of *transition rates*  $\Theta$  among

---

<sup>1</sup>The usual H3N1 strain was also present that year and hence the time series effectively combines two distinct outbreaks

the compartments. This SIR framework has been successfully used in modeling a range of infectious diseases, ranging from influenza [Merl et al 2009] to measles [Cauchemez and Ferguson 2008], and foot-and-mouth disease [Jewell et al 2009]. Probabilistically, this approach corresponds to imposing a Markovian structure at the group level and captures the intrinsic uncertainty through the stochasticity of the transitions taking place.

To model the observed strong seasonality of influenza outbreaks, we introduce a further level of uncertainty through a stochastic seasonal factor  $\{M_t\}$ . Thus, the transition rates between the compartmental classes are modulated by  $\{M_t\}$ , which can be interpreted as the seasonal presence of a new pathogen. This seasonal factor is evidently an abstract object, i.e. latent<sup>2</sup>. We note that its effect is indirect, since  $\{M_t\}$  only affects the *rates* of the transitions, not transitions themselves. We refer to LeStrat and Carrat [1999], Martínez-Beneito et al [2008] for related hidden Markov model representations of epidemics. There is further large literature on seasonal forcing in epidemics that focuses on deterministic SIR models using tools of dynamical systems [Keeling et al 2001, Dushoff et al 2004, Stone et al 2007] or multi-scale analysis [Kuske et al 2007].

Bayesian inference in compartmental models consists of estimating the transition rates between classes and the dynamic size of each class. This is by now a classical problem in biosurveillance; for example, a whole strand of literature is devoted to estimating the basic reproductive ratio  $\mathcal{R}$  which is the single most important parameter for predicting epidemic impact [Ball and Neal 2002, Cintron-Arias et al 2009, O'Neill 2002, Chowell et al 2009]. The complementary inference of latent states is treated in the aforementioned LeStrat and Carrat [1999], Martínez-Beneito et al [2008]. However, little research has been done for joint parameter and state estimation due to the associated computational challenges inherent in both Markov chain and sequential Monte Carlo approaches. In this paper we present a novel algorithm for such joint inference of  $\{M_t\}$  and outbreak parameters  $\Theta$  using sequential Monte Carlo.

Beyond pure inference, the ultimate objective of decision-makers is to mitigate the epidemic impact. This is achieved by implementing response policies including vaccination, quarantine, pharmaceutical or hospital treatment, information campaigns, etc. Since any policy involves budgetary or human resources, a balance is needed between costs due to epidemic morbidity and mortality and costs arising from policy actions. Moreover, to be optimal, a policy must be adaptive, i.e. rely on the latest collected information; thus, policy analysis is inherently linked to sequential inference. Quantitative approaches to such dynamic epidemic management include continuous-time Markov chain models [Merl et al 2009], Markov decision processes [Tanner et al 2008], systems of ordinary differential equations [Chowell et al 2009], agent-based representations [Halloran et al 2008] and stochastic control [Ludkovski and Niemi 2010]. However, existing methods are limited in adequately addressing the questions of unknown system dynamics, parameters and states. Below we demonstrate that regime-switching compartmental models in fact provide a flexible paradigm for analyzing these issues by allowing for accurate sequential inference.

## 2 General Setup

We consider a  $d$ -dimensional continuous-time jump-Markov process  $\{X_t\}$ ,  $X_t \in \mathbb{N}^d$  on a probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ , that evolves according to laws of a chemical reaction system. Namely,  $X_t^k$  denotes the (non-negative) number of species of class  $k = 1, \dots, d$  and there are  $\mathcal{Q}$  reaction *types* with corresponding stoichiometry vectors  $\Delta_q \in \mathbb{Z}^d$ ,  $q = 1, 2, \dots, \mathcal{Q}$ . The  $\Delta_q$ 's indicate the impact of a reaction on  $X_t$ : denoting by  $\tau_k$ ,  $k = 1, \dots$ , the reaction times, i.e. the jump times of  $\{X_t\}$ , if

---

<sup>2</sup>As stated by Grassly and Fraser [2006] “despite the near ubiquity of this phenomenon [seasonality], the causes and consequences of seasonal patterns of incidence are poorly understood”

the  $k$ -th reaction is of type  $q$ , then  $X_{\tau_k} = X_{\tau_{k-1}} + \Delta_q$ . Between reactions  $\{X_t\}$  is constant.

A convenient representation of  $\{X_t\}$  is via a multivariate marked point process  $\mathcal{X} \equiv (\tau_k, R_k)$ ,  $k = 1, 2, \dots$ , where

$$\tau_k := \inf\{t > \tau_{k-1} : X_t \neq X_{t-}\}, \quad k \geq 1, \quad \text{with } \tau_0 = 0,$$

are the transition epochs, and

$$R_k := X_{\tau_k} - X_{\tau_{k-1}} \in \{\Delta_1, \dots, \Delta_Q\}, \quad k \geq 1,$$

are the corresponding reactions which may be canonically identified with reaction types  $\{1, 2, \dots, Q\}$ . It is immediate that

$$X_t = X_0 + \sum_{\tau_k \leq t} R_k, \quad (1)$$

showing the equivalence of the two formulations. We also introduce the counting processes

$$N_t^q := \sum_{\tau_k \leq t} \mathbb{1}_{\{R_k = \Delta_q\}} \quad \text{and} \quad N_t^{q,i} := \sum_{\tau_k \leq t} \mathbb{1}_{\{R_k = \Delta_q, M_{\tau_k} = i\}}$$

for the number so far of each reaction by type, so that

$$X_t = X_0 + \sum_{q=1}^Q N_t^q \cdot \Delta_q.$$

The probability triple  $(\Omega, \mathcal{F}, \mathbb{P})$  also supports a finite-state jump process  $\{M_t\}$  that represents the modulating factor. Intuitively,  $M_t \in \{1, \dots, \mathcal{I}\}$  is a Markov chain with generator

$$G_M := (\mu_{ij}), \quad i, j \in \{1, \dots, \mathcal{I}\}. \quad (2)$$

Slightly more generally, we take  $G_M \equiv G_M(t, X_t)$ , allowing the transition rates of  $\{M_t\}$  to depend on time and the current state of  $\{X_t\}$ , so  $\{(X_t, M_t)\}$  is jointly Markov.

To complete the description of  $\{X_t\}$  and  $\{M_t\}$  we finally specify the transition rates  $\alpha$  of  $\{X_t\}$ , or equivalently the arrival rates of the counting process  $N^{q,i}$ . We assume that the corresponding propensity functions are of the form

$$\alpha_q(t, X_t, M_t) = \theta_q(M_t) \cdot h_q(t, X_t), \quad q = 1, \dots, Q,$$

where  $\Theta \equiv (\theta_1(1), \theta_1(2), \dots, \theta_Q(\mathcal{I}))$  are the reaction *rates*, modulated by  $\{M_t\}$ , and  $h$ 's are the mass action laws. Let

$$\bar{\alpha}(t, X_t, M_t) := \sum_{q=1}^Q \alpha_q(t, X_t, M_t) \quad (3)$$

be the total current arrival rate, and

$$f_q(t, X_t, M_t) := \frac{\alpha_q(t, X_t, M_t)}{\bar{\alpha}(t, X_t, M_t)} \quad (4)$$

be the conditional likelihood of the next reaction being of type  $q$ .

Due to the local Markovian structure, the distribution of the point process  $(\tau_k, R_k)$ , conditional on a path of  $\{M_t\}$  is given explicitly by [Amrein and Künsch 2012]

$$p(\tau_1, R_1, \dots, \tau_n, R_n | M_s, s \leq t) = \exp \left( - \int_0^t \bar{\alpha}(s, X_s, M_s) ds \right) \\ \times \prod_{k=1}^n \left( \sum_{q=1}^Q \alpha_q(\tau_k, X_{\tau_k-}, M_{\tau_k}) \cdot \mathbb{1}_{\{R_k=\Delta_q\}} \right), \quad (5)$$

where the first term accounts for the total arrival intensity on  $[0, t]$  and the second term for the likelihoods of the observed event types. Recall that if  $M_s$  is constant and  $\alpha$ 's are independent of  $t$ , then the inter-arrival times are exponentially distributed.

### 3 Inference

Statistical inference in the presented framework consists of estimating the reaction rates  $\Theta$  and the seasonal factor  $\{M_t\}$ . As already mentioned, throughout we assume that  $\{X_t\}$  is fully observed, i.e. by any date  $t$ , a full record of all reaction times  $\tau_k$  before  $t$  and corresponding reaction types  $R_k$  is available. To explain our method, we shall consider the following three cases:

- (a)  $\Theta$  unknown;  $\{M_t\}$  observed;
- (b)  $\Theta$  known;  $\{M_t\}$  unobserved;
- (c) Both  $\Theta$  unknown and  $\{M_t\}$  unobserved.

Case (c) is the main object of interest in our study; however to understand its properties, in the next sections we briefly review Cases (a) and (b).

Probabilistically, the three cases are distinguished by the different filtrations of observed information. Let  $(\mathcal{G}_t)$ ,

$$\mathcal{G}_t \equiv \mathcal{F}_t^{X,M,\Theta} := \sigma(X_s, M_s : 0 \leq s \leq t) \vee \sigma(\Theta), \quad (6)$$

denote the full filtration and using obvious notation let  $(\mathcal{F}_t^{X,M})$ ,  $(\mathcal{F}_t^{X,\Theta})$ , and  $(\mathcal{F}_t^X)$  denote the sub-filtrations corresponding to cases (a), (b), (c), respectively. Then our aim is to compute the (joint) posterior distribution  $\mathcal{Z}_t := p(M_t, \Theta | \mathcal{F}_t)$  of the seasonal factor and the parameters for the above choices of filtrations. For later use we also define the posterior probabilities  $\Pi_t := (\Pi_t^1, \dots, \Pi_t^I)$  where

$$\Pi_t^i = \mathbb{P}^\pi(M_t = i | \mathcal{F}_t^X), \quad i \in \{1, \dots, I\}, \quad (7)$$

where  $\mathbb{P}^\pi$  denotes the conditional probability measure given the prior distribution  $\Pi_0 = \pi$  of  $M_0$ .

#### 3.1 Conjugate Inference of Epidemic Parameters

If  $\{M_t\}$  is observed (i.e.  $(\mathcal{F}_t^{X,M})$  is available), then  $\{X_t\}$  forms a time-inhomogeneous Markov chain. In particular, a full description and analysis of  $\{X_t\}$  is possible on the intervals  $[\sigma_\ell, \sigma_{\ell+1}]$  where  $(\sigma_\ell)$  are the transition times of  $\{M_t\}$ :  $\sigma_{\ell+1} = \inf\{t \geq \sigma_\ell : M_t \neq M_{t-}\}$ .

It is well known that conjugate Bayesian updating of  $\Theta$  is available using Gamma priors [Boys et al 2008, Amrein and Künsch 2012]. Specifically, let

$$p(\Theta | \mathcal{F}_0^{X,M}) = \prod_{i=1}^{\mathcal{I}} \prod_{q=1}^{\mathcal{Q}} Ga(\theta_q(i); a_q^i, b_q^i),$$

where  $a.$  and  $b.$  denote the shape and rate parameters respectively of the Gamma distribution, be the independent priors of  $\theta_q$ . Then given  $\mathcal{F}_t^{X,M}$  we can express the *full data likelihood* as follows:

$$\begin{aligned} p(\Theta | \mathcal{F}_t^{X,M}) &= \prod_{i=1}^{\mathcal{I}} \prod_{q=1}^{\mathcal{Q}} Ga(\theta_q(i); a_q^i(t), b_q^i(t)), \\ a_q^i(t) &:= a_q^i + N_t^{q,i}, \\ b_q^i(t) &:= b_q^i + \int_0^t h_q(s, X_s) \mathbb{1}_{\{M_s=i\}} ds, \end{aligned} \quad (8)$$

where we recall that  $N_t^{q,i}$  is the number of observed reactions of type  $q$  during the  $i$ -regime.

We summarize those sufficient statistics as  $\mathbf{s}_t = (a_q^i(t), b_q^i(t))$ ,  $i = 1, \dots, \mathcal{I}$ ,  $q = 1, \dots, \mathcal{Q}$  and denote the updating by the function  $\mathcal{S}(\cdot)$  from (8), i.e.,

$$\mathbf{s}_T = \mathcal{S}(\mathbf{s}_t, t, T, M) := \left( a_q^i(t) + N_T^{q,i} - N_t^{q,i}, b_q^i(t) + \int_t^T h_q(s, X_s, M_s) \mathbb{1}_{\{M_s=i\}} ds \right)_{q=1}^{\mathcal{Q}}. \quad (9)$$

Relation (9) provides an explicit sequential way to update the posterior of  $\Theta$  along the trajectory of  $\{(X_t, M_t)\}$ .

### 3.2 Estimation of the Latent Seasonal Factor

In case (b), we assume that all rates  $\Theta$  are known and a full record of  $\{X_t\}$  is available, but the path of  $\{M_t\}$  is unobserved. Intuitively, inference of  $\{M_t\}$  is based on comparing the likelihoods of observed event epochs/types conditional on the possible values of the seasonal factor, using (5). Using the representation of  $\{X_t\}$  as a state-dependent marked doubly-stochastic Poisson process implies that  $\Pi_t$  can be described in closed-form and possesses piecewise-deterministic dynamics. Applying [Ludkovski and Sezer 2012, Prop 2.1] we obtain the following characterization

**Proposition 1.** *The sample paths of  $\{\Pi_t\}$  follow*

$$\begin{cases} \Pi_t = \vec{x}(t - \tau_k, \Pi_{\tau_k}, X_{\tau_k}), & \tau_k \leq t < \tau_{k+1}, k \in \mathbb{N} \\ \Pi_{\tau_k}^i = \frac{\alpha_{R_k}(\tau_k, X_{\tau_k-}, i) \Pi_{\tau_k-}^i}{\sum_{j=1}^{\mathcal{I}} \alpha_{R_k}(\tau_k, X_{\tau_k-}, j) \Pi_{\tau_k-}^j}, & i = 1, \dots, \mathcal{I}, \end{cases} \quad (10)$$

where the vector field  $\vec{x}(t, \vec{\pi}, X)$  is defined via

$$x^i(t, \vec{\pi}, X) = \frac{\mathbb{P}(\tau_1 > t, M_t = i | M_0 \sim \vec{\pi}, X_0 = X)}{\mathbb{P}(\tau_1 > t | M_0 \sim \vec{\pi}, X_0 = X)}, \quad (11)$$

$i \in \{1, \dots, \mathcal{I}\}$ , and has the explicit solution  $\mathbb{P}(\tau_1 > t, M_t = i | M_0 \sim \vec{\pi}, X_0 = X) = \vec{\pi} \cdot \exp(\int_0^t G_M(s, X) - A(s, X) ds)$  with  $A(s, X) := \text{diag}(\bar{\alpha}(s, X, \cdot))$  and

$$A(s, X)_{ii'} = \begin{cases} \bar{\alpha}(s, X, i) & \text{if } i = i'; \\ 0 & \text{otherwise.} \end{cases} \quad (12)$$

Proposition 1 completely identifies the distribution of the  $\mathcal{I}$ -dimensional posterior  $\Pi_t$  of  $M_t$  through the recursion (10).

### 3.3 Joint Sequential Inference using SMC

We now turn to our main case (c) where we only have access to  $(\mathcal{F}_t^X)$ , so that neither  $\{M_t\}$  nor  $\Theta$  is observed. In that case, even if all transitions of  $\{X_t\}$  are fully observed, the posterior distribution  $\mathcal{Z}_t = p(M_t, \Theta | \mathcal{F}_t^X)$  no longer admits any sufficient finite-dimensional statistics. Hence, no closed-form analysis is possible and we turn to constructing efficient numerical approximation schemes.

The sequential Monte Carlo approach to recursively update  $\mathcal{Z}_t$  consists of constructing a particle approximation

$$\mathcal{Z}_t \simeq \sum_{j=1}^J w_t^{(j)} \delta_{\{m_t^{(j)}, \theta^{(j)}\}}$$

where  $\delta$  is the Dirac delta function and each of the  $J$  particles is defined by its (normalized) weight  $w_t^{(j)} \geq 0$ , its  $M$ -location  $m_t^{(j)}$  and its parameter versions  $\theta^{(j)}$ . In other words,

$$\mathbb{P}((M_t, \Theta) \in A | \mathcal{F}_t^X) \simeq \sum_{(m_t^{(j)}, \theta^{(j)}) \in A} w_t^{(j)}.$$

The particles are updated using a propagation-selection scheme. However, since the parameters  $\Theta$  are fixed throughout, it is well known that naive SMC implementation typically leads to *particle degeneracy*, namely the diversity of  $\theta^{(j)}$ 's across the particles increasingly diminishes due to resampling. As we shall see in Section 4.2, this issue is acute for HMSKMs.

The availability of the sufficient statistics  $s_t$  from (9) for  $\Theta$  conditional on  $\{M_t\}$  allows to dramatically reduce degeneracy. We recall the early approach of Storvik [2002] who applied a basic bootstrap particle filter [Gordon et al 1993] on the pair  $(X_t, s_t)$ , where the likelihood of observations is evaluated by sampling from the posterior of  $\Theta$ , i.e.

$$p(X_t | M_t, s_t) \simeq p(X_t | M_t, \theta^{(j)}) \quad \theta^{(j)} \sim s_t \quad \text{i.i.d.} \quad (13)$$

### 3.4 Particle Learning in HMSKM

In our case, even more efficiency can be achieved through a resample-move SMC [Gilks and Berzuini 2001] which takes advantage of the explicit predictive likelihood of  $\mathcal{X}$  given  $\Theta$  (see (15) below) and leads to a version of the particle learning (PL) framework originally proposed within a discrete-time setting in Carvalho et al [2010, 2011]. The use of a *resample-move* algorithm allows direct sequential filtering of  $(M_t, s_t)$  rather than the static  $\Theta$ . Thus, the filtered distribution at time  $t$  is approximated by the particle cloud

$$\mathcal{Z}_t \simeq \mathcal{Z}_t^{(J)} := \sum_{j=1}^J w_t^{(j)} f_\Theta(\cdot | s_t^{(j)}) \delta_{\{m_t^{(j)}, s_t^{(j)}\}}, \quad (14)$$

where  $s_t^{(j)}$  are again the parameter sufficient statistics in (9). Thus, the conditional distribution of  $\Theta$ ,  $f_\Theta(\cdot | s_t)$ , is a product of independent Gamma's and the overall  $\mathcal{Z}_t^{(J)}$  represents the posterior of each  $\theta_q$  as a mixture of Gamma distributions.

```

Require: Priors  $\mathbf{s}_0$ ,  $\Pi_0$ , initial state  $X_0$ , number of particles  $J$ 
1: Sample  $m_0^{(j)} \sim \Pi_0$  i.i.d., set  $s_0^{(j)} \leftarrow \mathbf{s}_0$ ,  $j = 1, \dots, J$ 
2: loop { for  $k = 0, 1, \dots, J$ 
3:   Sample  $\theta^{(j)} \sim p(\theta | s_{\tau_k}^{(j)})$ ,  $j = 1, \dots, J$ 
4:   Calculate weights  $w_{k+1}^{(j)} \propto p(\tau_{k+1} - \tau_k, R_{k+1} | m_{\tau_k}^{(j)}, \theta^{(j)})$ 
5:   for  $j = 1, \dots, J$  do
6:     Re-sample  $j' \propto w_{k+1}^{(\cdot)}$  where  $j' \in \{1, \dots, J\}$ 
7:     Sample a trajectory  $m_{(\tau_k, \tau_{k+1})}^{(j')}$  using the conditional law  $p(M_{(\tau_k, \tau_{k+1})} | m_{\tau_k}^{(j')}, \theta^{(j')}, \tau_{k+1}, R_{k+1})$ 
8:     Update  $s_{k+1}^{(j)} \leftarrow \mathcal{S}(s_k^{(j')}, \tau_k, \tau_{k+1}, m_{(\tau_k, \tau_{k+1})}^{(j)})$ 
9:   end for
10: end loop

```

Algorithm 1: Particle Learning for Hidden Markov Stochastic Kinetic Model

Algorithm 1 summarizes in pseudo-code the steps of the proposed particle learning algorithm. Its main steps are computing particle weights in (4) for resampling using the predictive likelihood of the next event, forward propagation step (6) using the conditional law of the environment factor, and updating of the sufficient statistics step (7) for the parameters. Overall, besides the analytical results detailed below, only the ability to simulate  $\{M_t\}$  and  $\{X_t\}$  is needed to implement the above Monte Carlo scheme, highlighting the flexibility of PL. A basic simulation method for SKM that is exact and can always be used is the Gillespie algorithm (see e.g. Wilkinson [2006]), here slightly extended to take into account additional transition times of  $\{M_t\}$ .

To calculate the *predictive likelihood* of the next inter-arrival interval  $\tau_{k+1} - \tau_k$  and reaction type  $R_{k+1}$  conditional on  $M_{\tau_k}$  and parameters  $\Theta$  we rely on the analytic expression in (5),

$$p(\tau_{k+1} - \tau_k, R_{k+1} | M_{\tau_k}, \Theta) = \sum_{i=1}^{\mathcal{I}} p(R_{k+1} | M_{\tau_{k+1}} = i, \Theta) \times p(\tau_{k+1} - \tau_k, M_{\tau_{k+1}} = i | M_{\tau_k}, \Theta). \quad (15)$$

The first term on the right-hand-side is  $f_r(\tau_{k+1}, X_{\tau_{k+1}-}, M_{\tau_{k+1}})$  using the parameters  $\Theta$ , and the second term is

$$\mathbb{P}(\tau_{k+1} - \tau_k = t, M_{\tau_{k+1}} = i | M_{\tau_k} = i', \Theta) := P_{i'i}(t) \bar{\alpha}(\tau_{k+1}, X_{\tau_{k+1}-}, i) \quad (16)$$

with  $P_{i'i}(t)$  being an element of the matrix exponential  $P(t) = e^{\int_0^t G_M(s, X) - A(s, X) ds}$ , see (12). When there are just two latent states,  $|\mathcal{I}| = 2$ ,  $P(t)$  can be computed explicitly using eigenvector decomposition.

The *conditional law* of  $\{M_t\}$  given  $\tau_{k+1} - \tau_k, R_{k+1}$  is not available in closed form. However, using Bayes rule

$$\begin{aligned} & p(M_{(\tau_k, \tau_{k+1})} | M_{\tau_k}, s_{\tau_k}, \Theta, \tau_{k+1}, R_{k+1}) \\ & \propto p(\tau_{k+1} - \tau_k, R_{k+1} | M_{(\tau_k, \tau_{k+1})}, \Theta) p(M_{(\tau_k, \tau_{k+1})} | M_{\tau_k}) \end{aligned} \quad (17)$$

where

$$\begin{aligned} & p(\tau_{k+1} - \tau_k, R_{k+1} | M_{(\tau_k, \tau_{k+1})}, \Theta) \\ & = \exp \left\{ - \int_{\tau_k}^{\tau_{k+1}} \bar{\alpha}(s, X_{\tau_k}, M_s) ds \right\} \cdot \alpha_{R_{k+1}}(\tau_{k+1}, X_{\tau_k}, M_{\tau_{k+1}}) \end{aligned} \quad (18)$$

and  $p(M_{(\tau_k, \tau_{k+1}]} | M_{\tau_k})$  is determined from the transition matrix of  $\{M_t\}$ . To implement (17) we use a rejection sampling step relying on the fact that there is an easy upper bound of (18) :

$$\begin{aligned} p(\tau_{k+1} - \tau_k, R_{k+1} | M_{(\tau_k, \tau_{k+1}]}, \Theta) \\ \leq \exp\left\{(\tau_k - \tau_{k+1}) \min_{i,s} \bar{\alpha}(s, X_{\tau_k}, i)\right\} \cdot \max_i \alpha_{R_{k+1}}(\tau_{k+1}, X_{\tau_k}, i). \end{aligned} \quad (19)$$

Thus, we simulate using the unconditional law  $p(M_{(\tau_k, \tau_{k+1}]} | M_{\tau_k})$  and then accept the simulation with probability given by the ratio between (18) and (19). Since  $\tau_k$  are typically tightly spaced, the above conditional likelihoods are all close to each other, requiring only a few additional simulations (acceptance probability is usually  $> 95\%$ ).

Finally, as already mentioned, the sufficient statistics for  $\Theta$  are always conjugate-Gamma with the updating given explicitly in (9). We note that for typographical convenience in Algorithm 1 resampling takes place at each reaction time  $\tau_k$ . In practice, any other resampling frequency can be chosen; in that case the weights are updated accordingly until resampling takes place. Also, a variety of resampling schemes (multinomial, residual, stratified, etc.) are available [Cappé et al 2005] and can be used to lower Monte Carlo error.

Like all SMC methods, PL still exhibits sample impoverishment which implies that the filtering error  $\|\mathcal{Z}_T^{(J)} - \mathcal{Z}_T\|$  (in an appropriate metric) grows exponentially in  $T$ . Thus, exponentially more particles are needed to control the Monte Carlo error in terms of the number of observations. We therefore recommend utilizing PL on a fixed horizon  $T$  as in our application below.

## 4 Model of Seasonal Epidemics

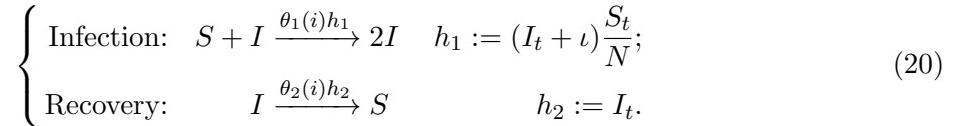
We now return to our main example of a Markov-modulated chemical reaction system — a compartmental model of seasonally-forced endemic diseases. As a simple example, we shall analyze a classical stochastic SIR-type model (20) of epidemics that incorporates a latent seasonal factor  $\{M_t\}$ . For concreteness, we phrase our discussion in terms of the human influenza virus.

A basic description of an endemic disease such as flu can be provided using an SIS compartmental model [Andersson and Britton 2000], which features just two population compartments of Susceptibles  $\{S_t\}$  and Infecteds  $\{I_t\}$ . Since only partial immunity is available against influenza (it is common for an individual, especially children, to have several flu episodes in one season), the Recovered compartment is omitted and we assume that upon recovery individuals immediately pass back into the susceptible pool. Other models of endemic diseases are reviewed in Nåsell [2002].

Let  $M_t \in \{1, 2\}$  denote the seasonal factor at date  $t$ , with  $M_t = 1$  representing low flu season and  $M_t = 2$  high flu season. We assume a closed population of constant size  $N := S_t + I_t$  that represents a fixed geographic area (such as a college campus, a town, or a county). We assume that  $\{M_t\}$  forms a time-stationary Markov chain with infinitesimal generator

$$G_M := \begin{pmatrix} -\mu_{12} & \mu_{12} \\ \mu_{21} & -\mu_{21} \end{pmatrix}.$$

In other words, holding times for  $\{M_t\}$  in regime  $i$  are exponentially distributed with mean  $\mu_{ij}^{-1}$ . Conditional on  $M_t = i$ ,  $\{X_t\} \equiv \{S_t, I_t\}$  is a jump-Markov process involving two reaction types



Thus, new infections take place according to the law of mass action [Andersson and Britton 2000], where the infection rate is driven by the possible pairings between infected and susceptible individuals (assuming homogenous mixing of the full population). Additionally, we add an “immigration of infecteds” rate  $\iota$  which can be viewed as an external reservoir of the flu (e.g. from travellers) that provides a constant source of additional infections. We use this term to prevent a stochastic fade-out of the epidemic and guarantee endemicity. There are two epidemic parameters  $\Theta \equiv (\theta_1, \theta_2)$ , interpreted as infectiousness ( $\theta_1$ ) and mean recovery time ( $1/\theta_2$ ). Since  $S_t = N - I_t$ ,  $\{I_t\}$  summarizes the epidemic state and we omit  $S_t$  from further discussion.

For simplicity, we assume that seasonal variations affect only the contact rate  $\theta_1(M_t)$ , so that  $\theta_2(M_t) = \theta_2$  is constant. We moreover assume that the effect on the contact rate is multiplicative,

$$\theta_1(2) = (1 + SF)\theta_1(1),$$

for a known seasonality impact ratio  $SF > 0$ . This is meant to model the case where the seasonality increases probability of susceptibles becoming infected (due e.g. to cold weather) but has no impact on the severity of the flu once infected.

As described above, given  $M$ , the key state  $\{I_t\}$  forms a recurrent Markov chain on  $\{0, 1, \dots, N\}$ . Using the notation in Section 2, we have  $R_k \in \{-1, 1\}$  with

$$\alpha_1(I, i) := \theta_1(i) \frac{(I + \iota)(N - I)}{N}, \quad \alpha_{-1}(I, i) := \theta_2 I, \quad (21)$$

for  $i = 1, 2$ , and

$$f_r(I, i) = \begin{cases} \frac{\theta_1(i)(I+\iota)(N-I)/N}{\theta_1(i)(I+\iota)(N-I)/N + \theta_2 I} & \text{if } r = 1; \\ \frac{\theta_2 I}{\theta_1(i)(I+\iota)(N-I)/N + \theta_2 I} & \text{if } r = -1. \end{cases} \quad (22)$$

Since  $\{M_t\}$  only takes on two values, its posterior is described by the one-dimensional probability process

$$\Pi_t^2 := \mathbb{P}(M_t = 2 | \mathcal{F}_t^I).$$

Applying Proposition 1 and using the fact that  $\theta_2$  is independent of  $\{M_t\}$  leads to the following simple dynamics of  $\Pi_t^2$ :

**Corollary 1.** *The evolution of  $\{\Pi_t^2\}$  is characterized on each  $t \in [\tau_k, \tau_{k+1})$  by*

$$\frac{d\Pi_t^2}{dt} = (1 - 2\Pi_t^2)\mu_{21} - SF \cdot \theta_1 \frac{(N - I_{\tau_k})(I_{\tau_k} + \iota)}{N} \Pi_t^2(1 - \Pi_t^2), \quad (23)$$

$$\text{and } \Pi_{\tau_k}^2 = \begin{cases} \frac{(1+SF)\Pi_{\tau_k-}^2}{1+SF\Pi_{\tau_k-}^2} > \Pi_{\tau_k-}^2, & \text{if } R_k = 1, \\ \Pi_{\tau_k-}^2, & \text{if } R_k = -1. \end{cases} \quad (24)$$

An explicit solution to (23) can be obtained by computing the eigen-pairs of the corresponding matrix  $A(I)$ .

The particle learning Algorithm 1 can be straightforwardly applied in this model using particles of the form  $(m_t^{(j)}, s_t^{(j)}) \in \{1, 2\} \times \mathbb{R}_+^4$ . Since  $\theta_2$  is not modulated, its posterior is in fact *deterministic* given  $\mathcal{F}_t^I$ , i.e. the same for all particles and only the two sufficient statistics of  $\theta_1$  need to be

recorded:

$$\begin{aligned}
p\left(\theta_1 | \mathcal{F}_t^I, \{m_s^{(j)}, 0 \leq s \leq t\}\right) &= Ga(a_1 + N_t^1, b_1(t)); \\
p\left(\theta_2 | \mathcal{F}_t^I\right) &= Ga(a_2 + N_t^2, b_2 + \int_0^t I_s ds); \\
b_1(t) &:= b_1 + \int_0^t \left(1 + SF \mathbb{1}_{\{m_s^{(j)} = 2\}}\right) \frac{(I_s + \iota)(N - I_s)}{N} ds,
\end{aligned}$$

where  $\mathbf{s}_0 = (a_1, a_2, b_1, b_2)$  is the original Gamma prior of  $\Theta$ .

#### 4.1 Illustration

To illustrate the above algorithms, this section presents several numerical experiments. Table 1 summarizes the parameters used. We consider a single flu season lasting 9 months (approximately September through May), which includes on average a 6-month long high-season (so that  $\mu_{21}^{-1} = 365/2$ ). Such a model is meant to capture a single seasonal cycle as is commonly done by US epidemiological agencies on a Fall-Spring basis, see Figure 1.

We consider a fixed population of  $N = 10,000$  individuals with an initial  $I_0 = 50$  infecteds. During the low season,  $\theta_1(1) < \theta_2$  meaning that the epidemic would on its own fade out; the disease remains endemic through contact with outside infecteds, with a carrying capacity (i.e. the long-run expected level of Infecteds freezing environmental fluctuations) of about  $\lim_{t \rightarrow \infty} \mathbb{E}[I_t | M_s = 1 \forall s] \simeq 50$ . In the high season,  $\theta_1(2) > \theta_2$  so that an outbreak begins. Even though the infectiousness rate increases by just  $SF = 15\%$ , the resulting carrying capacity jumps to over 500, i.e. more than 5% of total. This illustrates that even small changes in the contact rate can have a dramatic effect on the equilibrium disease incidence. In line with common estimates, we use average infectiousness period of about 4 days,  $\theta_2 \simeq 0.25$ .

Table 1: Parameter values used. All rates are daily.

Parameter	Meaning	Value
$\mu_{12}$	Transition rate to high season	6/365
$\mu_{21}$	Transition rate to low season	2/365
$\theta_2$	Recovery rate	0.25
$\theta_1(1)$	Low-season infectiousness	0.235
$\theta_1(2)$	High-season infectiousness	0.27025
$SF$	Seasonality effect	0.15
$\iota$	Immigration of outside infecteds	2
$N$	Population size	10,000
$T$	Time horizon (days)	273
$(M_0, I_0)$	Initial Condition	(0,50)
$(a_1(0), b_1(0))$	Initial Priors for $\theta_1$	(25,100)
$(a_2(0), b_2(0))$	Initial Priors for $\theta_2$	(25,100)

Figure 2 shows a sample trajectory of the infected population count over the nine months in conjunction with the underlying seasonal factor  $\{M_t\}$ . In this scenario, starting with low-season, high season  $\{M_t = 2\}$  begins on day 61 and ends on day 182, lasting just over four months (compared to average high season length of 1/2-year). The seasonality effect is clear, as soon after the beginning of the high season,  $\{I_t\}$  begins an upward trend which is reversed once  $M_t = 1$

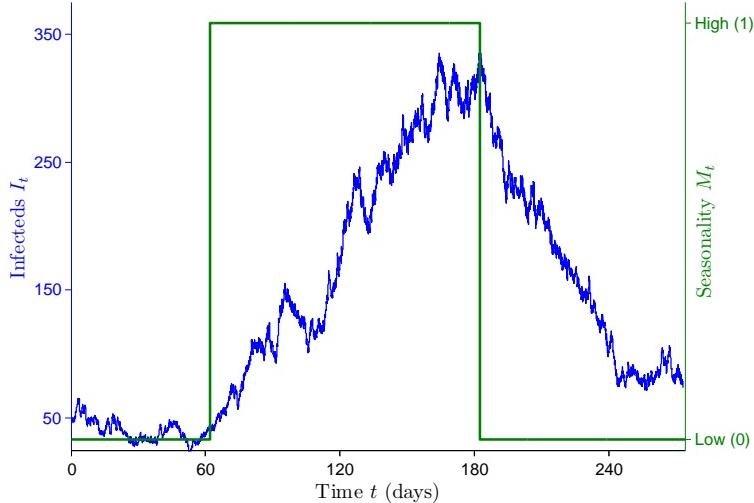


Figure 2: Sample path of  $\{M_t, I_t\}$  over a course of 9 months. There are a total of 20079 transitions.

again. Nevertheless, we observe a lot of stochastic fluctuations against these main trends (e.g. a significant decrease in  $I_t$  around day 100). Overall, there were a total of 10051 infections recorded over the period, corresponding to roughly each member of the population becoming infected once, in line with observed statistics on influenza. Peak number of infecteds was 335 on day 164.

Starting with a rather vague prior for  $\Theta$ , Figure 3 shows the sequential Bayesian inference of  $\Theta$  assuming that the trajectory of  $\{M_t, I_t\}$  shown in Figure 2 is fully observed. We note that the posterior means apparently converge to the true values and the posterior credibility interval (CI) narrows at roughly a hyperbolic rate over time. As data is accumulated, the oscillations in the posterior distributions decrease quickly.

Figure 4 presents the results from the other nested model (b), namely assuming known parameters  $\Theta$  but unobserved seasonal factor  $\{M_t\}$ . In Figure 4 we show the posterior probability  $\Pi_t^2$  of the high season over the same trajectory of  $\{I_t\}$  shown in Figure 2. We note that generally the filter is able to well-identify the present seasonality effect and responds very quickly when  $\{M_t\}$  changes (see the very sharp drop around  $t = 185$ ). The dynamic lag between change in the true  $\{M_t\}$  and the response by the filter is on the order of 10-20 days which is quite fast and would be difficult to identify with a “naked eye” looking at the trajectory of  $\{I_t\}$ . At the same time, the filter  $\Pi_t^2$  is highly sensitive to the local behavior of  $\{I_t\}$  making it very noisy. For example the mentioned drop in infecteds around  $t = 100$  causes posterior likelihood of high season to decrease from over 95% to as low as 35%, albeit with a sharp reversal once the upward trend is re-established. The underlying piecewise-deterministic behavior of  $\{\Pi_t^2\}$  (see Proposition 1) is highlighted in the inset figure which clearly shows the discrete upward jumps of  $\{\Pi_t^2\}$  when new infections are recorded.

Finally, Figure 5 shows the output from running the SMC algorithm for joint inference of  $(M_t, \Theta)$  using  $J = 5000$  particles. Uncertainty about both parameters and seasonal factor makes the posterior credible intervals wider compared to Figure 3. In terms of the posterior probability of the high season, the resulting  $\Pi_t^2$  is less volatile compared to Figure 4 and responds slightly slower to underlying regime shifts.

The implemented instance of Algorithm 1 is somewhat computationally intensive since it re-

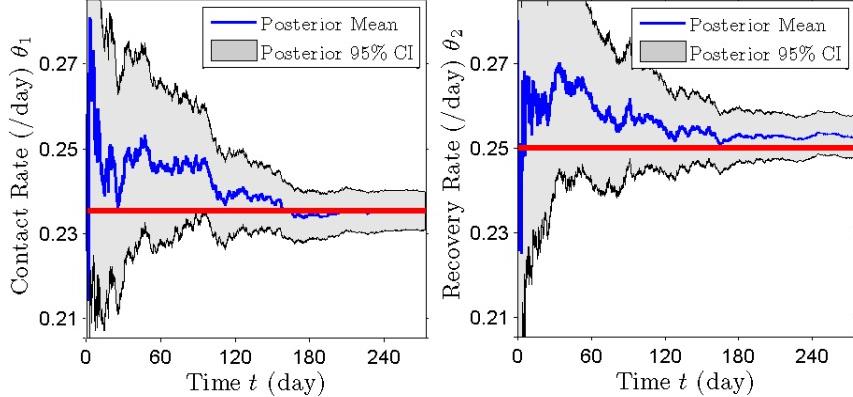


Figure 3: Posterior mean and 95% credibility interval of  $\Theta \equiv (\theta_1, \theta_2)$  over the sample path of  $\{M_t, I_t\}$  in Figure 2. Solid horizontal lines indicate the true values used.

quires repeated simulation of paths of  $\{M_t\}$  and evaluation of the predictive and conditional likelihoods over each interval  $[\tau_k, \tau_{k+1})$  (so over 20,000 times in the Figures shown) and for each particle ( $J = 5000$ ). However, we note that all these computations are *exact*, so the only noise present is from Monte Carlo re-sampling. Running time to generate Figure 5 is about two minutes on a typical 2011 laptop.

#### 4.2 Comparison to Other Inference Methods

From the time-series estimation point of view, the stochastic systems we consider are characterized by long time-series and short inter-event periods. It is well known that the basic challenge of Sequential Monte Carlo over long horizons is particle degeneracy since repeated re-sampling necessarily throws some information away and cuts down on diversity of the empirical particle cloud. Degeneracy is exacerbated when one is required to estimate constant parameters. Without further steps, the basic bootstrap filter of Gordon et al [1993] will degenerate, almost surely as  $t \rightarrow \infty$ , to a point mass estimate of the posterior density. A popular simple solution is the Liu and West [2001] algorithm (henceforth LW) that introduces adjustment moves to particle versions  $\theta^{(j)}$  of  $\Theta$ . We implemented LW as a comparison to the presented PL algorithm and found that degeneracy is still prevalent over the second half of the season even with as many as  $J = 5000$  particles. As such, use of the sufficient statistics  $s_t$  for  $\Theta$ -posteriors is *crucial* for HSMKMs.

We also implemented the Storvik [2002] filter. Compared to the PL filter, its main difference is that Storvik [2002] applies propagate-resample steps, and as such does not require analytic form of the predictive likelihood (17) in Step 4 of Algorithm 1. Instead, one simply uses

$$w_{\tau_k}^{(j)} \propto \exp \left( - \int_{\tau_{k-1}}^{\tau_k} \bar{\alpha}(I_s, m_s^{(j)}) ds \right) \cdot \alpha_{R_k}^{(j)}(I_{\tau_{k-1}}, m_{\tau_k}^{(j)}),$$

where the propensity rates  $\alpha$ 's use the sampled particle versions  $\theta^{(j)} \sim s_{\tau_{k-1}}$ . The Storvik algorithm runs slightly faster since it does not need to evaluate the matrix exponential in (23).

To quantitatively compare the performance of the described SMC algorithms, we started by creating a benchmark using a PL filter with  $2 \cdot 10^4$  particles. We then generated 100 runs of the PL, Storvik and LW schemes each using  $J = 2000$  particles and residual resampling for

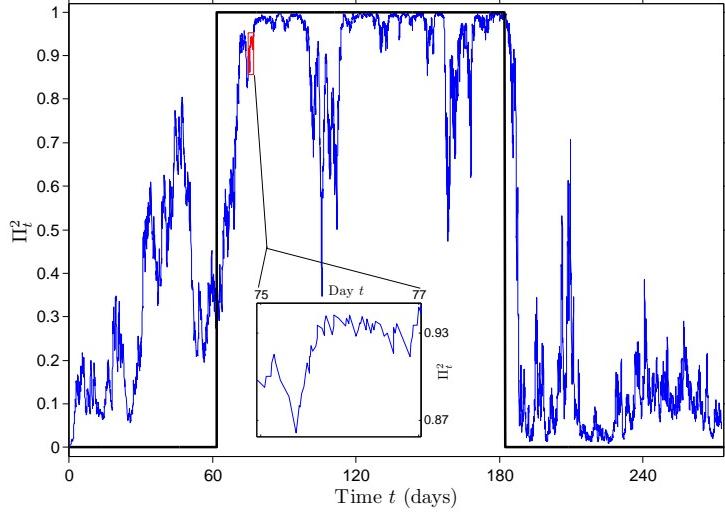


Figure 4: Posterior probability  $\Pi_t^2 = \mathbb{P}(M_t = 2|\mathcal{F}_t^I)$  over the sample path of  $\{I_t\}$  in Figure 2, assuming known  $(\theta_1, \theta_2)$ . For comparison we also show the respective true trajectory of  $\{M_t\}$ .

the realization of  $\{M_t, I_t\}$  in Figure 2. Because in the case study the second reaction rate  $\theta_2$  is independent of  $\{M_t\}$ , its sufficient statistic  $s_t^2$  is the same for all particles and is computed exactly by both Storvik and PL algorithms for any filter size  $J$ . On the other hand, the estimate of  $\theta_1$  is highly sensitive to correct tracking of  $\{M_t\}$  over time. We first compare the 95% coverage probabilities with respect to the true  $\theta_1$ , i.e. how frequently does the true value  $\theta_1 = 0.235$  belong to the corresponding 95% posterior CI obtained from SMC. The results for the two time-points of  $T_1 = 120$  and  $T_2 = 270$  days are summarized in Table 2. We find that the PL algorithm performs best; the Storvik algorithm has somewhat higher MC errors but still maintains particle diversity. The LW algorithm posteriors start to collapse mid-way through the data and completely degenerate by the end, so that their coverage probabilities are nil. Figure 6 further shows that the posterior 95% CI of the PL algorithm includes fewer outliers, in other words nearly all runs of the algorithm recover the correct posterior.

In Table 2 we also compare the standard error  $\mathbb{E} [\|\Pi_t^2 - \hat{\Pi}_t^2\|_2]$  of the posterior probability of the high season  $\Pi_t^2$  with respect to the benchmark filter  $\hat{\Pi}_t^2$ . We observe that compared to PL, the Storvik scheme is more prone to ‘‘losing track’’ of  $M_t$ , in other words the Monte Carlo runs appear to be more leptocurtic. The LW algorithm in fact usually tracks  $\hat{\Pi}_t^2$  well but also had several completely failed runs and clearly suffers from particle degeneracy. Overall, this analysis confirms our preference for PL; its advantages can be compared to the improvement from the bootstrap to the auxiliary particle filter in the classical SMC setup.

### 4.3 Discussion

The presented model is clearly stylized and would not be able to capture all the features of real epidemics. Since both  $\Theta$  and  $\{M_t\}$  are assumed unobserved, it may still be possible to fit real data even under such model mis-specification. Nevertheless, in this section we discuss some of adjustments that may be made for achieving further realism.

While year-to-year epidemics arise in different times, there are clear patterns which imply that

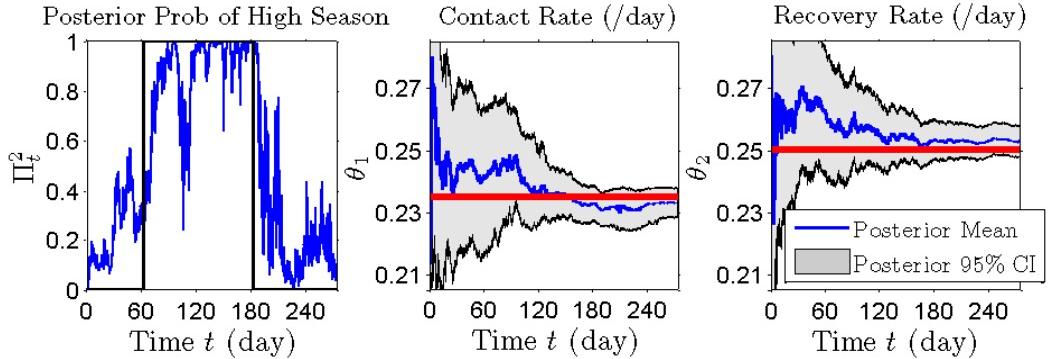


Figure 5: Joint inference of  $(M_t, \Theta)$  over the trajectory of  $\{I_t\}$  in Figure 2. Left panel shows the posterior probability of the high season  $\Pi_t^2 = \mathbb{P}(M_t = 2 | \mathcal{F}_t^I)$ . The other two panels show the posterior median and 95% credibility interval of the two parameters  $\theta_1, \theta_2$ . The SMC algorithm used  $J = 5000$  particles. All other parameters are from Table 1.

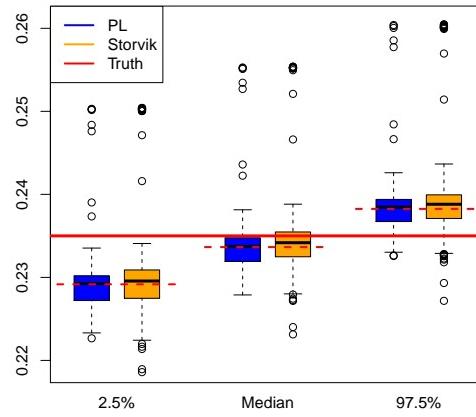


Figure 6: Posterior quantiles of  $\theta_1$  at  $T_2 = 270$  days using the PL and Storvik algorithms with  $J = 2000$  particles. The histograms show the median and 95% CI quantiles across 100 runs of each algorithm. The true value of  $\theta_1$  is indicated by the solid horizontal line, the dashed lines indicate the corresponding quantiles from a benchmark run of PL with 20000 particles.

Table 2: Comparison of SMC algorithm performance. All algorithms used  $J = 2000$  particles and we took  $T_1 = 120$ ,  $T_2 = 270$  days for the path of  $\{M_t, I_t\}$  in Figure 2. The LW algorithm had a tuning parameter of  $h = 0.97$ .

	PL	Storvik	LW
95% Cov. Prob of $\theta_1$ at $T_1$	0.96	0.87	0.32
95% Cov. Prob of $\theta_1$ at $T_2$	0.74	0.65	0.01
Std Error of $\Pi_t^2$ at $T_1$	0.200	0.272	0.455
Std Error of $\Pi_t^2$ at $T_2$	0.056	0.105	0.037

assuming time-stationary transition rates of  $\{M_t\}$  is not reasonable. Our general setup allows for  $G_M(t)$  in (2) to be time-dependent and could be used to capture these patterns. Other calendar-year effects, such as the impact of the school-year start [He et al 2010], can be deterministically added. A more serious concern is the assumption of known seasonal transition rates  $\mu_{ij}$ ; however in our 2-state example those may be reasonably well-calibrated using historical data such as that used in Figure 1. Our numerical experiments also show that mis-specification of  $\mu_{ij}$  is not a serious problem for sequential inference of  $\Pi_t$ .

The model (20) assumes that there is a single infectiousness parameter  $\theta_1$  which remains constant throughout, such that the actual contact rates are of the form  $\alpha_1(I_t, M_t) = \theta_1(1 + SF\mathbb{1}_{\{M_t=2\}})h_1(I_t)$ . This representation is for convenience only; it is straightforward to consider the case where we separately carry along priors for each  $\theta_1(i)$ . We could also consider the case where each new season leads to a “fresh”  $\theta_1$  (e.g. from a new strain of the pathogen), i.e. the contact rate on each interval  $[\sigma_\ell, \sigma_{\ell+1})$ , where  $(\sigma_\ell)$  are the transition epochs of  $\{M_t\}$  is  $\theta_{1,(\ell)}(M_{\sigma_\ell}) \sim p(\theta)$ . In that case, the sufficient statistics  $s_t^{(j)}$  are simply reset to the original prior when the corresponding particle copy of  $m_t^{(j)}$  changes states. However, the difficulty with this method is that  $\theta_1(i)$ ’s are not independent. For instance, to have the interpretation of  $M_t = 2$  being the high season, we require  $\theta_1(2) > \theta_1(1)$ . This precludes the simple specification of the respective marginals through a Gamma prior. The assumption of a constant seasonal factor  $SF$  is a convenient work-around which enforces the epidemiological meaning of seasonality.

Our choice of a two-compartment SIS model was due to its simplicity. Since our methods can handle any HMSKM, it is immediate to extend to further compartments (e.g. an SEIRS model with  $X_t = (S_t, E_t, I_t, R_t)$  that would further include the Exposed compartment of individuals who are infected but not infectious, and the Recovered compartment for individuals who have temporary immunity) or multiple seasonal regimes (e.g. adding a third “pandemic” regime to capture outbreaks like 2009 H1N1 influenza). By modifying the propensity functions  $h_q(t, X_t)$  one can also refine the modeling of population mixing [He et al 2010], incorporate external immigration of individuals, or include age-structured populations.

Finally, an extra possibility is to allow two-way feedback between the epidemic state  $\{X_t\}$  and the seasonal factor  $\{M_t\}$ . Thus, beyond the modulation of the transition times of  $X$  by  $M$ , we can also introduce effect of  $X$  on  $M$  through  $G_M = G_M(t, X_t)$ . In other words, the seasonality dynamics are themselves affected by the epidemic. For instance, a large outbreak could be due to a long-surviving pathogen which in turn prolongs the expected length of a high-season. Alternatively, one could consider the case where each infection increases the chances of a genetic mutation of the pathogen, thereby decreasing  $\mu_{21}$ . We refer to Ludkovski [2012a] for SMC algorithms to address this possibility.

The most severe constraint of our model is the assumption of fully observing  $\{I_t\}$ . In principle this could be achieved by exhaustive monitoring of all individuals’ status. More realistically,

we view this model as an idealized case of biosurveillance which contends with extrinsic model uncertainty regarding  $\Theta$  and  $\{M_t\}$  while eschewing missing data. This then provides a useful benchmark to analyze data collection quality. We refer to Ludkovski and Niemi [2011] for a related setup with discrete-time observations based on binomial sub-sampling which also admits a PL algorithm.

## 5 Optimized Policy Response

As mentioned in the introduction, an important use of sequential inference is for policy response. In this section we briefly investigate such biosurveillance decision-making using the developed inference tools.

Public health policy makers act sequentially as outbreak data is collected. Their aim is to balance total costs which consist of morbidity costs  $C_M$  associated with the pathogen, and costs  $C_A$  associated with policy actions. Denote by  $\phi_t \in \mathbb{R}$  the policy implemented at date  $t$ , where we encode the action space as a subset of the real line. We consider time-additive costs on a given horizon  $[0, T]$  of the form

$$\begin{aligned} C_M &:= \int_0^T c(I_t) dt \quad \text{and} \\ C_A &:= \int_0^T \phi_t dt + \sum_{\phi_t \neq \phi_{t-}} K(\Delta\phi_t), \end{aligned} \tag{25}$$

where  $c(I_t)$  are the morbidity costs expressed as instantaneous rates, and without loss of generality we take  $\phi_t$  to be the instantaneous cost of the respective action. The last term in (25) with  $\Delta\phi_t \equiv \phi_t - \phi_{t-}$  corresponds to potential additional *start-up* costs  $K(\cdot)$  associated with changing a policy.

In contrast to the original presentation, now  $\phi_t$  dynamically drives the evolution of  $\{I_t, M_t\}$ . The impact of policy actions can be either direct or indirect. Direct actions influence the transition rates, whereby  $\Theta = \Theta(\phi_t)$ . For example, quarantine can cut down infectiousness rates in a population. Indirect actions influence the transition matrix  $G_M = G_M(\phi_t)$  of  $\{M_t\}$ , e.g. making low-season regime  $M_t = 1$  more likely. These can be interpreted as prophylactic measures that mitigate the seasonal effect and contain the disease at its baseline morbidity. In either case, since applied policies must be based only on currently available information, the control  $\{\phi_t\}$  is required to be  $\mathcal{F}_t^I$ -adapted.

We assume that the overall objective is to minimize average (expected) costs over the interval  $[0, T]$  across all potential dynamic policies,

$$\inf_{(\phi_t)} \mathbb{E}^\phi [C_M + \lambda C_A] = \inf_\phi \mathbb{E}^\phi \left[ \int_0^T c(I_t) + \lambda \{\phi_t + K(\Delta\phi_t)\} dt \right], \tag{26}$$

where  $\lambda$  is a Lagrange multiplier and to emphasize the impact of  $\phi$  we denote the resulting probability measure as  $\mathbb{P}^\phi$ . As  $\lambda$  is increased, the budget constraints tighten; for  $\lambda = 0$  the optimization is solely about minimizing expected morbidity.

The optimization problem (26) is a partially observed stochastic control problem. The presence of unobserved quantities makes it nonstandard; a general approach is to reduce it to a standard setting by passing from  $I_t$  to the augmented hyperstate  $\mathcal{Z}_t$ . As we have seen before,  $\{\mathcal{Z}_t\}$  is finite dimensional if at least one of  $\{M_t\}$  or  $\Theta$  is observed, but is infinite-dimensional in the main case of interest (c) requiring joint inference. Using the Markov structure, the optimal policy response

$\phi_t^*$  at time  $t$  is a function of the posterior state  $\mathcal{Z}_t$ ,  $\phi_t^* = \Phi(I_t, \mathcal{Z}_t, \phi_{t-})$ , for some *strategy rule*  $\Phi$ . Analytic treatment of such infinite-dimensional control problems is generally intractable; see Ludkovski and Niemi [2010], Ludkovski [2012b] for flexible numerical approximations that also rely on SMC.

Rather than carry out a full optimization, we investigate below some simple heuristics for such rules  $\Phi$ . As our case study we consider a simple example with a binary policy response,  $\phi_t \in \{0, 1\}$  with  $\phi_t = 1$  indicating the implementation of preventive measures at date  $t$ . We assume that transition rates  $\Theta$  are fixed but policy makers can influence the environment  $\{M_t\}$  through

$$G_M(t) \Big|_{\phi_t=0} = \begin{pmatrix} -6 & 6 \\ 2 & -2 \end{pmatrix} \quad \text{and} \quad G_M(t) \Big|_{\phi_t=1} = \begin{pmatrix} -1 & 1 \\ 8 & -8 \end{pmatrix}.$$

Thus, when counter-measures are enacted,  $\{M_t\}$  is much more likely to be in the low state. Indeed, while without any actions the high season is expected to last 6 months, with mitigation it will only last an average of  $52/8 \simeq 6$  weeks.

With an on/off decision-making, the dynamic policy rule is described through its action regions  $D_0$  and  $D_1$ , such that a response is initiated as soon as  $(I_t, \mathcal{Z}_t) \in D_1$  and  $\phi_{t-} = 0$ , and terminated as soon as  $(I_t, \mathcal{Z}_t) \in D_0$  and  $\phi_{t-} = 0$ . Note that we expect to have a hysteresis region where the existing policy, whatever it may be, is continued, since there is no sufficiently strong evidence to change the response. To compare, we consider three potential classes of policies that are distinguished by the information used:

**The infecteds-based control** relies directly on  $I_t$ . Since start (resp. end) of outbreaks is characterized by persistent upward (resp. downward) trend, we take  $D_1^{Infid} = \{I_t - \mathcal{I}_t^{(\kappa)} > \bar{I}\}$ ,  $D_0^{Infid} = \{I_t - \mathcal{I}_t^{(\kappa)} < \underline{I}\}$ , where

$$\mathcal{I}_t^{(\kappa)} := \frac{1}{t} \int_0^t e^{-\kappa s} I_s ds$$

is an exponentially weighted moving average of the number of infecteds using discount weight  $\kappa$ . Thus, measures are initiated when  $I_t$  is sufficiently above its moving average (strong upward trend), and stopped when  $I_t - \mathcal{I}_t^{(\kappa)}$  is sufficiently negative. From our experiments, using a 14-day moving average offers a good way of capturing trends. This is a simple policy that requires no inference and can be seen as rule of thumb to provide response when  $\{I_t\}$  is growing.

**The Bayesian policy** incorporates the full history of observations and the prior beliefs through the posterior distribution  $\mathcal{Z}_t$ . As a simple choice we consider only the posterior of  $\{M_t\}$  and analyze  $D_1^{Bayes} = \{\Pi_t^2 \geq \bar{\pi}\}$  and  $D_0^{Bayes} = \{\Pi_t^2 \leq \underline{\pi}\}$ . Thus, countermeasures are started once the posterior probability of  $\{M_t = 2\}$  is above  $\bar{\pi}$  (overwhelming evidence of a high season), and are stopped once the posterior probability is below  $\underline{\pi}$ .

**The Oracle policy** takes  $D_0^{Orcl} = \{M_t = 1\}$ ,  $D_1^{Orcl} = \{M_t = 2\}$ , or  $\phi_t = \mathbb{1}_{\{M_t=2\}}$ , i.e. the control is applied precisely during the high season. This is an idealized benchmark since it assumes that the policy-maker actually observes the latent seasonality variable. We note that it is still not the globally optimal benchmark since in principle it may not be worthwhile to respond *immediately* as soon as the high season begins (e.g. if  $I_t$  is still low then ostensibly  $M_t$  could revert back to its low state quickly without the need for costly interventions).

Clearly, the above list is not exhaustive and is for illustration purposes only. An infinite variety of further rules can be constructed. For example, including posterior information about  $\Theta$  is clearly

relevant to understand the severity of the outbreak and its likely future course. In general, an automated way to build an optimal policy is more appropriate than heuristics, albeit at a cost of losing some of the simplicity and intuition for the decision maker.

Once a class of policies is chosen, the decision rule needs to be optimized by minimizing over the described parametric forms such as  $(\underline{\pi}, \bar{\pi})$  or  $(\underline{I}, \bar{I})$ . Note that even for a fixed rule, the expected costs are not analytically available since the distribution of  $I_t$  or  $M_t$  under  $\mathbb{P}^\phi$  is not explicitly computable. We therefore resort to predictive analysis via Monte Carlo. Namely, fixing the policy rule, we simulate a large number (several hundred in our example below) of scenarios, i.e. trajectories of  $(I_t, M_t, \phi_t)$ , and then average the resulting scenario costs to approximate  $\mathbb{E}^\phi[C_M + \lambda C_A]$ .

Simulation under the controlled measure  $\mathbb{P}^\phi$  is straightforward thanks to the strong Markov properties of  $\{\mathcal{Z}_t\}$  and is summarized in Algorithm 2. For simplicity we restrict to the case where  $\phi_t$  only changes at event times  $\tau_k$ . For these simulations we draw the actual parameters  $\Theta$  independently from the given prior  $p(\Theta)$ , i.e. assuming the model is correctly specified and then generate  $\{I_t, M_t\}$  using the Gillespie algorithm.

```
Require:  $(M_0, I_0, \mathcal{Z}_0, \phi_0)$ 
1: Sample outbreak parameters  $\Theta \sim \mathcal{Z}_0$ 
2:  $s \leftarrow 0$ 
3: loop
4:   Simulate the next regime change date  $\sigma \geq s$  of  $\{M_t\}$  using the generator  $G_M(\phi_s)$ 
5:   Simulate the next transition  $\tau \geq s$  of  $\{I_t\}$  conditional on  $(M_s, \phi_s, \Theta)$ 
6:    $\rho \leftarrow \sigma \wedge \tau$ 
7:   Save  $M_{[s, \rho]}$  and  $I_{[s, \rho]}$ 
8:   Using PL Algorithm 1 update the filter  $\{\mathcal{Z}_t\}$  on  $[s, \rho]$ 
9:   Update the policy  $\phi_\rho \leftarrow \Phi(I_\rho, \mathcal{Z}_\rho, \phi_s)$ 
10:  Set  $s \leftarrow \rho$ 
11: end loop
```

Algorithm 2: Simulation of a controlled epidemic model

## 5.1 Cost Functionals

There are many possible summary statistics to evaluate the relative merit of different mitigation strategies. Among morbidity measures, one can consider average number of infecteds  $\mathbb{E}^\phi[\frac{1}{T} \int_0^T I_t dt]$ , maximum infecteds  $\mathbb{E}^\phi[\max_{0 \leq t \leq T} I_t]$ , or the proportion of time that  $I_t$  is above some level  $I_{high}$ ,  $\mathbb{E}^\phi[\int_0^T \mathbb{1}_{\{I_t \geq I_{high}\}} dt]$ . One can also include metrics regarding the response, such as the average length of time countermeasures are enacted  $\mathbb{E}^\phi[\int_0^T \mathbb{1}_{\{\phi_t=1\}} dt]$ , the number of times counter-measures are started  $\mathbb{E}^\phi[\sum_{t \leq T} \mathbb{1}_{\{\Delta\phi_t \neq 0\}}]$ , etc. To illustrate, we consider the following two examples of cost functionals:

$$c_1(I, \phi) := \int_0^T I_t + 0.02(I_t - 200)_+^2 + 50 \cdot \mathbb{1}_{\{\phi_t=1\}} dt; \quad (27a)$$

$$c_2(I, \phi) := \int_0^T (I_t + 1000 \cdot \mathbb{1}_{\{I_t > 300\}} + 200 \cdot \mathbb{1}_{\{\phi_t=1\}}) dt + 1400 \cdot \sum_{t \leq T} \mathbb{1}_{\{\Delta\phi_t \neq 0\}}. \quad (27b)$$

The cost functional  $c_1$  has a piecewise-quadratic cost in terms of the number of infecteds  $I_t$ . Here a “soft” threshold of  $I_{high} = 200$  infecteds is applied, as well as a basic morbidity cost

that is proportional to  $I_t$ . This  $c_1$  also mildly penalizes the amount of time counter-measures are applied through the third term in (27a). The cost functional  $c_2$  is geared more towards minimizing mitigation resources. It has much higher policy costs of 200 per day and also rewards policy stability (i.e. avoiding too many changes in policy or “chattering”). The latter is taken into account in (27b) through the penalty  $K(\Delta\phi_t) = 1400 \cdot \mathbb{1}_{\{\Delta\phi_t=1\}}$  which imposes a start-up cost of 1400 (week’s worth of policy costs) each time the counter-measures are begun. In terms of outbreak costs,  $c_2$  considers total number of infecteds-days and additionally imposes a discontinuous penalty whenever  $I_t$  exceeds 300 (i.e. 3% of the population) which can be thought of as a “hard” (but high) target regarding tolerable number of infecteds.

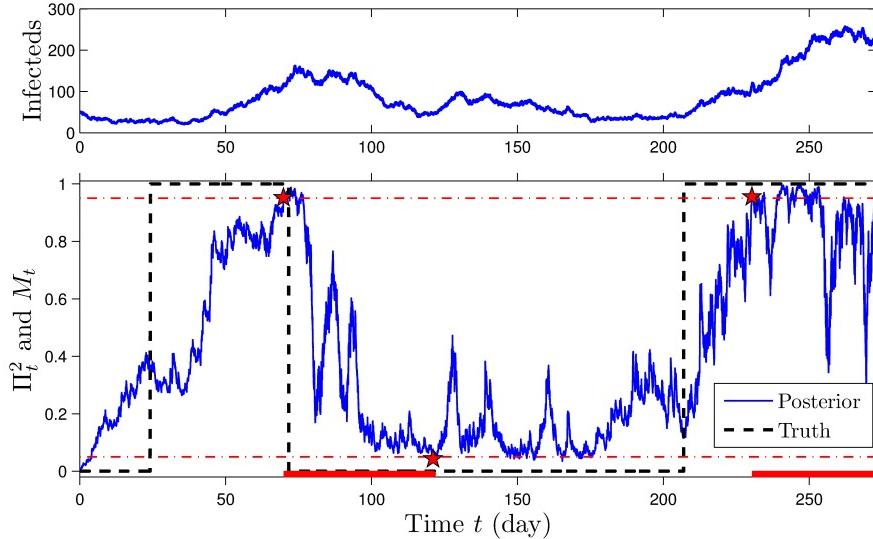


Figure 7: Dynamic Bayesian control of  $\{M_t\}$  over a sample trajectory of the outbreak. The top panel shows the infected numbers  $\{I_t\}$ ; the bottom panel shows the true seasonal factor  $\{M_t\}$  and the posterior probability of high season  $\{\Pi_t^2\}$ . Countermeasures are applied as soon as  $\Pi_t^2 > 0.95 = \bar{\pi}$  and stopped once  $\Pi_t^2 < 0.05 = \underline{\pi}$ . The intervals of action are indicated by the solid bars on the  $x$ -axis, and the times of policy changes are marked with stars. The PL algorithm used  $J = 5000$  particles. All other parameters are from Table 1.

To generate outbreaks, we fix  $\theta_2 = 0.25$  and sample  $\theta_1 \sim Ga(1700, 20 \cdot 365)$ , which roughly means  $\theta_1 \in [0.21, 0.25]$ . The remaining parameters, including  $SF = 0.15$ , are from Table 1. Figure 7 shows the resulting Bayesian-type dynamic policy on a sample trajectory of  $\{I_t\}$ . In this example, counter-measures start once the probability of being in the high season is at least 95% and are ended once  $\Pi_t^2 < 0.05$  and  $I_t < I_{high} = 200$ ; the latter is to make sure that the outbreak is fully contained. In the figure, the filter reacts rather slowly to the first outbreak, perhaps due to lower than normal numbers of infecteds at its onset. Coincidentally, once counter-measures are finally started,  $\{M_t\}$  reverts back to low season almost immediately. The second outbreak begins around seven months and is responded to within 20 days. At the end of the simulation,  $M_t = 2$  remains in the high season. Note that the unrealistic assumption (solely for simplicity of presentation) that  $\{M_t\}$  goes back to its original transition matrix once  $\phi_t = 0$  (even after counter-measures were applied previously) implies that multiple high seasons are likely to occur over the nine months.

Policy Type	Thresholds	Ave. $I_t$ ( $\pm 3$ )	Policy Freq. ( $\pm 1\%$ )	Ave. Days $I_t > 300$ ( $\pm 1.5$ )	Ave. Policy Start-Ups ( $\pm 0.1$ )	$\mathbb{E}^\phi[c_1]$ ('000s) ( $\pm 4$ )	$\mathbb{E}^\phi[c_2]$ ('000s) ( $\pm 2.5$ )
Baseline	(Do-Nothing)	174	0%	54	0	187	102
Oracle		95	41%	15	4.8	60	70
Infecteds	(20,-10)	92	37%	13	2.6	50	66
Infecteds	(40,-20)	106	28%	17	5.3	59	65
Bayesian	(0.80,0.01)	91	49%	12	2.7	48	67
Bayesian	(0.95,0.05)	108	30%	19	2.6	56	64

Table 3: Expected costs and summary statistics of selected response strategies over 500 simulated scenarios. The chosen Infecteds-based rule is  $D_1^{Infid} = \{I_t - \mathcal{I}_t^{(14)} > \bar{I}\}$ ,  $D_0^{Infid} = \{I_t - \mathcal{I}_t^{(\kappa)} < \underline{I} \cap I_t < I_{high}\}$  and the chosen Bayesian rule is  $D_1^{Bayes} = \{\Pi_t^2 \geq \bar{\pi}\}$  and  $D_0^{Bayes} = \{\Pi_t^2 \leq \underline{\pi} \cap I_t < I_{high}\}$  for the specified thresholds. Bayesian algorithms used  $J = 3000$  particles. The Baseline policy is to do nothing  $\phi_t \equiv 0$ , and the Oracle policy is  $\phi_t = \mathbb{1}_{\{M_t=2\}}$ . Standard deviations of the computed expected values are in brackets.

The performance of several mitigation policies is compared in Table 3. The table demonstrates that depending on the priorities of the policy makers, different mitigation strategies should be considered. Crucially, one must consider the trade-off between fast response and potentially unnecessary interventions. This trade-off is clearly observed with some strategies being more aggressive (and hence producing lower expected morbidity) and other strategies being more conservative. For instance, the Bayesian policy with  $(\underline{\pi}, \bar{\pi}) = (0.01, 0.8)$  is much more aggressive than a similar policy with  $(\underline{\pi}, \bar{\pi}) = (0.05, 0.95)$ , as it starts counter-measures as soon as there is at least 80% chance of being in the high season, and continues them until the posterior probability drops below 1% (compared to starting at  $\Pi_t^2 > 95\%$  and stopping as soon as  $\Pi_t^2 < 5\%$ ). Not surprisingly, it reduces average number of infecteds by nearly 15% and the average number of days when there are more than 300 infecteds by over 30% in comparison. This comes at a cost of over 40 additional days on average when countermeasures are applied. Which policy is better therefore depends on the weightings (cf. Lagrange multiplier  $\lambda$  in (26)) placed on the different ingredients of the cost functional; here  $\mathbb{E}^\phi[c_1]$  is smaller for  $(\underline{\pi}, \bar{\pi}) = (0.01, 0.8)$ , while  $\mathbb{E}^\phi[c_2]$  is smaller for  $(\underline{\pi}, \bar{\pi}) = (0.05, 0.95)$ .

Overall, we find that strategies based on Bayesian inference seem able to outperform policies based on  $I_t$  only, though the difference is not very statistically significant. Of course, this result must be tempered as we did not perform an exhaustive optimization and it obviously depends on the cost functionals considered. We also note that strategies differ a lot in achieving similar results. For instance the Bayesian policy with  $(\underline{\pi}, \bar{\pi}) = (0.01, 0.8)$  has similar expected costs using the functional  $c_1$  as the Infecteds-based policy with  $(\underline{I}, \bar{I}) = (-10, 20)$ . However, the latter is much more variable (over 5.3 expected policy start-ups) in time, though lasting shorter periods (20% less frequently).

As a final comparison, Table 3 points out that no action at all is clearly sub-optimal and whatever the aims of the policy makers, *some* mitigation is obviously beneficial. Moreover, the idealized Oracle control is in fact not optimal either, since it does not account for the dynamics of  $\{I_t\}$  and by tracking  $\{M_t\}$  exactly tends to act/end too quickly. Further tailoring and refinement of risk metrics is obviously needed for practical use and raises a host of interesting inter-disciplinary questions that will be explored in upcoming works.

## References

- Amrein M, Künsch H (2012) Rate estimation in partially observed Markov jump processes with measurement errors. *Stat. Comp.* 22(2):513–526
- Andersson H, Britton T (2000) Stochastic epidemic models and their statistical analysis, Lecture Notes in Statistics, vol 151. Springer, New York
- Ball F, Neal P (2002) A general model for stochastic SIR epidemics with two levels of mixing. *Math Biosci* 180:73–102
- Bellomy A (2011) Influenza surveillance activities local update report. Tech. rep., Santa Barbara County Public Health Department, [http://www.countyofsb.org/phd/epi.aspx?id=23610&ekmensex=15074a7f\\_1152\\_1272\\_23610\\_1](http://www.countyofsb.org/phd/epi.aspx?id=23610&ekmensex=15074a7f_1152_1272_23610_1)
- Boys RJ, Wilkinson DJ, Kirkwood TB (2008) Bayesian inference for a discretely observed stochastic kinetic model. *Stat. Comp.* 18(2):125–135
- Cappé O, Moulines E, Rydén T (2005) Inference in hidden Markov models. Springer Series in Statistics, Springer, New York
- Carvalho CM, Johannes M, Lopes HF, Polson N (2010) Particle learning and smoothing. *Stat. Sci.* 25:88–106
- Carvalho CM, Johannes M, Lopes HF, Polson N (2011) Particle learning for sequential Bayesian computation. *Bayes. Stat.* 9:317–360
- Cauchemez S, Ferguson N (2008) Likelihood-based estimation of continuous-time epidemic models from time-series data: application to measles transmission in london. *J. R. Soc. Interface* 5(25):885–897
- Chowell G, Viboud C, Wang X, Bertozzi S, Miller M (2009) Adaptive vaccination strategies to mitigate pandemic influenza: Mexico as a case study. *PLoS One* 4(12):e8164
- Cintron-Arias A, Castillo-Chavez C, Bettencourt L, Lloyd A, Banks H (2009) The estimation of the effective reproductive number from disease outbreak data. *Math. Biosci. Eng.* 6 (2):261–282
- Doucet A, de Freitas N, Gordon N (eds) (2001) Sequential Monte Carlo methods in practice. Statistics for Engineering and Information Science, Springer, New York
- Dukic V, Lopes H, Polson N (2010) Tracking flu epidemic using Google Flu Trends and particle learning. Available at <http://www.ssrn.com/abstract=1513705>
- Dushoff J, Plotkin J, Levin S, Earn D (2004) Dynamical resonance can account for seasonality of influenza epidemics. *Proc. Nat. Acad. Sci. USA* 101(48):16,915
- Gilks WR, Berzuini C (2001) Following a moving target: Monte Carlo inference for dynamic Bayesian models. *J. Roy. Stat. Soc. B* 63:127–146
- Golightly A, Wilkinson D (2011) Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface Focus* 1(6):807–820
- Golightly A, Wilkinson DJ (2006) Bayesian sequential inference for stochastic kinetic biochemical network models. *J. Comp. Biol.* 13(3):838–851

- Gordon NJ, Salmond DJ, Smith AFM (1993) Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc. F* 140:107–113
- Grassly N, Fraser C (2006) Seasonal infectious disease epidemiology. *Proc. Roy. Soc. B - Biol. Sci.* 273(1600):2541–2550
- Halloran ME, Ferguson NM, Eubank S, Ira M Longini J, Cummings DAT, Lewis B, Xu S, Fraser C, Vullikanti A, Germann TC, Wagener D, Beckman R, Kadau K, Barrett C, Macken CA, Burke DS, Cooley P (2008) Modeling targeted layered containment of an influenza pandemic in the united states. *Proc. Nat. Acad. Sci. USA* 105 (12):4639–4644
- He D, Ionides E, King A (2010) Plug-and-play inference for disease dynamics: measles in large and small populations as a case study. *J. Roy. Soc. Interface* 7(43):271–283
- Jewell C, Kypraios T, Neal P, Roberts G (2009) Bayesian analysis for emerging infectious diseases. *Bayes. Anal.* 4(3):465–496
- Keeling M, Rohani P, Grenfell B (2001) Seasonally forced disease dynamics explored as switching between attractors. *Physica D* 148(3-4):317–335
- Kuske R, Gordillo L, Greenwood P (2007) Sustained oscillations via coherence resonance in SIR. *J. Theor. Biol.* 245(3):459–469
- Lawson A (2009) Bayesian disease mapping: hierarchical modeling in spatial epidemiology, vol 20. Chapman & Hall/CRC, New York
- LeStrat Y, Carrat F (1999) Monitoring epidemiologic surveillance data using hidden Markov models. *Stat. Med.* 18:3463–3478
- Liu J, West M (2001) Combined parameter and state estimation in simulation-based filtering. In: Sequential Monte Carlo methods in practice, Stat. Eng. Inf. Sci., Springer, New York, pp 197–223
- Ludkovski M (2012a) Bayesian quickest detection with observation-changepoint feedback. In: Proceedings of the 2012 Conference on Decision and Control, Maui HI Dec 9-12, 2012
- Ludkovski M (2012b) Monte Carlo methods for adaptive disorder problems. In: Carmona R, Moral PD, Hu P, Oudjane N (eds) Numerical Methods in Finance, Springer Proceedings in Mathematics, vol 12, Springer, pp 83–112
- Ludkovski M, Niemi J (2010) Optimal dynamic policies for influenza management. *Stat. Comm. Inf. Diseases* 2(1):article 5
- Ludkovski M, Niemi J (2011) Optimal disease outbreak decisions with noisy, delayed observations. Preprint
- Ludkovski M, Sezer S (2012) Finite horizon decision timing with partially observable Poisson processes. *Stoch. Models* 28(2):207–247
- Martínez-Beneito C, López-Quílez A, López-Maside A (2008) Bayesian Markov switching models for the early detection of influenza epidemics. *Stat. Med.* 27:4455–4468
- Merl D, Johnson R, Gramacy B, Mangel M (2009) A statistical framework for the adaptive management of epidemiological interventions. *PLoS ONE* 4(6):e5087

- Nåsell I (2002) Stochastic models of some endemic infections. *Math. Biosci.* 179(1):1–19
- Niemi J (2009) Bayesian analysis and computational methods for dynamic modeling. PhD thesis, Duke Univ
- O'Neill PD (2002) A tutorial introduction to Bayesian inference for stochastic epidemic models using Markov chain Monte Carlo methods. *Math. Biosci.* 180:103–114
- Stone L, Olinky R, Huppert A (2007) Seasonal dynamics of recurrent epidemics. *Nature* 446(7135):533–536
- Storvik G (2002) Particle filters in state space models with the presence of unknown static parameters. *IEEE Tran. Signal Proces.* 50(2):281–289
- Tanner MW, Sattenspiel L, Ntiamo L (2008) Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Math. Biosci.* 215(2):144–151
- Wilkinson DJ (2006) Stochastic Modelling for Systems Biology. Chapman & Hall/CRC, London